AI artificial intelligence big model multi-model multi-modal general agent research and development actual combat [illegible]2025v1.4 e-book.

●[illegible]- [illegible] - [illegible] - [illegible] - [illegible] - [illegible]- [illegible] - [illegible] - [illegible]Transformer [illegible] - [illegible] Q-learning[illegible] - [illegible]- [illegible] - [illegible] - [illegible] - [illegible]- [illegible] - [illegible] Transformer [illegible] - [illegible] - [illegible]- [illegible] - [illegible] - [illegible]- [illegible] - [illegible] Python[illegible]C++[illegible] - [illegible] PyTorch[illegible]TensorFlow [illegible] Hugging Face Transformers[illegible]DeepSpeed[illegible]Megatron-LM [illegible] - [illegible] Git[illegible]- [illegible] - [illegible] - [illegible] Horovod[illegible]NCCL [illegible] GPU[illegible] - [illegible]- [illegible] - [illegible] - [illegible]- [illegible] - [illegible] - [illegible]- [illegible] - [illegible] - [illegible]- [illegible] - [illegible] - [illegible]

●[illegible]- [illegible] Transformer [illegible]- [illegible]- [illegible]- [illegible]- [illegible] GPU[illegible]TPU [illegible]- [illegible]- [illegible]- [illegible]- [illegible] Python[illegible]C++[illegible] PyTorch[illegible]TensorFlow [illegible]- [illegible]

●[illegible]---### **[illegible]**1. **[illegible]** - **[illegible]**[illegible] Transformer[illegible]MoE[illegible]CoT[illegible] Qwen2.5-Omni [illegible] TMRoPE [illegible] RAG[illegible] - **[illegible]**[illegible]PEFT[illegible]RLHF[illegible]

□□□□ - **□□□□□□**□□□□□□□□□□□□□□/□□□□□□□□□□□ OpenAI □ GPT-4o □□□□□□□□□□□2. **□□□□□□□** - **□□□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□ Qwen2.5-Omni □ Thinker-Talker □□□□□□□□□□□□□□□□□□□□□□ - **□□□□□□**□□□□□□□□□□□□□ TMRoPE□□□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□ OCR□□□□□□□□□□□□□□3. **□□□□Agent□□□□□** - **□□□□□□**□□□ MetaGPT□□□□□□□□□□□□AutoGen□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□□**□□□□□□□□□□□□□ Multi-Agent □□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□“□□□□□”□“□□□□”□□□---### **□□□□□□□□□**1. **□□□□□□** - **□□□□□□□□**□□ - **□□□□□□**□□□□□□□□□ Open6DOR□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□□□□**□□□2. **□□□□□□□□** - **□□□□□□□□**□□□□“□□□+□□□+□□□”□□□□□□□□□□□□□□□□□□ RAG □□□□□□□□□□□ Agent □□□ - **□□□□□□**□□□□□□□□□□□□□□□□□ MoE □□□□□□□□□□□□RTX 3090□□□□□H100 □□□□□□□□□□ - **□□□□□□□□**□□□□□□□□□□□□□ Qwen2.5-Omni □ 300ms □□□□□□□□□□□□□□□□□□□□□□□□3. **□□□□□□** - **□□□□□□**□□□ 95%□ - **□□□□□□**□□□□□□□ RAG□□□□ Memo RAG□□□□ KV □□□□□□□□□□□□□□□□ - **□□□□□□□□**□□□□□□□□□□□□□ RLHF□□□□ AI □□□□□□□□□□□□ AI □□□□□□□□□□□□□---### **□□□□□□□□□□□□□**1. **□□□□□□□□** - **□□□□□□**□□ - **□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□2. **□□□□□□** - **□□□□□□**□□□“□□□□”□□□□□□□□□□□□□□□□□ - **□□□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□“□□□□”□□□□□□---### **□□□□□□□□□**1. **□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□Agent □□□□□□□□ 2. **□□□□□□**□□□□□□□□□□□□□□ HuggingGPT□□□□□□□□□□□□□□□□□□□□□ 3. **□□□□□□**□□□□□□□□□□□□□□ OpenAI □ Figure AI □□□□□□□□□□□□□□□□□□□---### **□□□□□□**- **□□□□□□**□□□□□□□□□□□TurboS □□□□□□□□□□□□□□□□□□□□ Agent □□□□□□□□□□□□□ - **□□□□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□**□□□“AI+□□”□□□□□---**□□**□□□□□□□□□□□□□□□□□□□“□□□□+□□□□+□□□□”□□□□□□□□□□□□□□□□□□□□□□Agent □□

●□□□- □□□□□□□□ - □□□□ □□ - □□□□□□□□ □□□ - □□□□□□ □□- □□□□□□□□ - □□□□ □□□ - □□□□ □□Transformer □□□□□□□□□□□□□□□□□□□□ - □□□□ □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ Q-learning□□□□□□□□□□□ - □□□□□□- □□□□□□□□□□□□□ - □□□□□□□□□□□□□□ □□ - □□□□□□□ □□□- □□□□□□ - □□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ Transformer □□□□□□□□□□□□□□□□□□□□□□□□□□□ - □□□□□□□□ □□□ - □□□□□□□□□ □□- □□□□ - □□□□□□□□□ □□ - □□□□□□□ □□- □□□□□□□□□□ - □□□□□□□□□ □□□ Python□C++□□□ - □□□□□□□□□ □□□□□□ PyTorch□TensorFlow □□□□□□□□□□□□□□□□□□□□□□□ Hugging Face Transformers□DeepSpeed□Megatron-LM □□□□□□□□□□□□□□□ - □□□□□□□□□ □□□□□□□□□□□□□ Git□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□□□□□□□□□□ - □□□□□□□ □□□ - □□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□ Horovod□NCCL □□□□□□□ GPU□□□□□□□□□□□□□□□□□□ -

[illegible] [illegible]- [illegible] - [illegible] [illegible] - [illegible] [illegible]- [illegible] - [illegible] [illegible] - [illegible] [illegible]- [illegible] - [illegible] [illegible] - [illegible] [illegible]- [illegible] - [illegible] [illegible] - [illegible] [illegible]

●[illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]

●### [illegible]---#### **[illegible]**1. **[illegible]** - **[illegible]**[illegible]56.5%[illegible]AMD [illegible]AI [illegible]2025 [illegible]AI [illegible]40%[illegible] - **[illegible]**[illegible]8 [illegible]OpenAI [illegible]GPT-4o [illegible]Mistral [illegible]Le Chat[illegible] - **[illegible]**[illegible]“[illegible]2025”[illegible]SSG [illegible]610 [illegible]82%[illegible]2. **[illegible]** - **[illegible]**[illegible]Meta [illegible]“[illegible]”[illegible] - **[illegible]**[illegible]“AI+”[illegible]2025 [illegible]217 [illegible]IP[illegible]“AI [illegible]”[illegible]60 [illegible] - **[illegible]**[illegible]

□□□□□□□□□□□□1090 □□□□□ AI □□□□□□ Cyber Valley □□□□□□---#### **□□□□□□□□□□□□□□□**1. **□□□□□□**- **□□**□□□ AI □□□□□□9010 □□□□□□□□□□□□□□□□□□□□□□□□□□□□6.5 □□□□ - **□□**□□□□□□□□□□413 □□□□□□□□□□□□□□□□□□□70 □□□□BAT □□□□□□□□□2025 □□□□□ 55 □ AI □□□□□ - **□□**□□□□□□□□□□□□□□□ 30%□□□□□□□□□□□□□□□□□□□□□□□□□ AI □□□□□□□□□□□□2. **□□□□□□**- **□□**□□□□□MIT □□□□□□□□□□□□□□□□OpenAI □□□□□□□□□□□□□□□□□ - **□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□ DeepSeek-R1 □□□□□□□□□□□□□□□ 1/10□□□□□□□□□□ - **□□**□□□“□□ 4.0”□□□□---#### **□□□□□□□□□□□□□**1. **□□□□□□□□□**- **2025 □□□**□□□ AI □□□□□□ 1186 □□□□□□□□□□□□□□□□□ 45%□□□□□□□□□□□□□□□□□□□ - **□□□□□**□ - **□□□□**□□□□□□□□□20 □□□L4 □□□□□□□□□□120 □□□□□□□□□□ FSD □□□□□□□□□□□□ - **□□□□**□□□□□□□□□□□□□15 □□□2030 □□□□□ Walker X □□□□□□□□□□ - **AI □□□**□□□ SSG □□□□□□ 21.1%□□□□□□□ 82%□GPU □□□□□□ 13 □□2. **□□□□□□**- **□□**□□□□□□□□□□□□□□□□□□2025 □ AI+□□□□□□□ 500 □□□□□□□□□□□ 9.3 □□□□□□ - **□□**□□□ B □□□□□□□□□□□□□□□□□□□□□□□□□□□□□ 60%□OpenAI □□ API □□□□□□□ 50 □□□□ - **□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□ AI□□□□ Le Chat □□□□□□□□□□□□□□□□□□□□□---#### **□□□□□□□□□□□□□□□□□**1. **□□□□□□□□**- **□□□□**□AI+□□□□□□□□□□□□□ 30%□□□□□□□□□□□□□□□ 2030 □□ 676 □□□□□□□□□ 1/3□ - **□□□□□**□AI □□□□□□□ 15%□□□□□□□□□□□ 25%□□□□□□□□□□□□□14nm □□□□□2. **□□□□□□**- **2030 □□□ AI □□□□**□□ 15 □□□□□□□□□□ 30%□□□ 40%□□□ 20%□ - **□□□□□**□□□□□□□□□□□□□□ 60%□□□□□□□□□□□□□□□□ 25%□□AI □□□□□□□ 10%□□□ 40%□□---### **□□** □□□□□□□□□□□□□□□□□□□□□□“□□□□□□□□□□□□□□□□”□□

●●□□ □□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□□□ Python □□□□□□ if x = 5: □□□□□□ if x == 5: □□□□□□□□□□□□□ □□□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□□□ Python □□□□□□□□□□□□□□□□□□□□□□□□ def calculate_sum(): □□□□□ def CalculateSum(): □□□□□□□ □□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□ □□□□□□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□□□ Java □□□□□□□□ int □□□□□ age □□□□□□□□□□□□□ age = "25" □□□□□□□□□□□□□□□□□ □□□□□□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□ □□□□□□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- □□□512□□□□□□□ 256□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ □□□□□□ - □□- □□□ PyTorch □□ □□□□□□□□ - □□□- □□

●□□□□□□□□□□□□□□□□□□□□□□□□□□ □□□□□□ - □□□- □□□□□□□□□□□□□□□□□□□□□□□□□□□□□Adagrad □□□□□□□□□□□□□□□□□□□□□□□□□□□ □□□□ - □□□- □□ □□□□□□□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ train_data_loader □□□□□□□□□□□- □□□□□□□□□

[illegible] - [illegible]- [illegible] - GPU/CPU [illegible] GPU [illegible] CPU [illegible]- [illegible] - [illegible]- [illegible]

●●[illegible] - [illegible]- [illegible] AES [illegible] - [illegible]- [illegible] - [illegible]- [illegible] - [illegible] SQL [illegible]- [illegible] - [illegible]- [illegible] Python [illegible]

●[illegible] 1. [illegible]```pythondef load_model(): return torch.load("/home/user/fixed_path/model.pth")``` [illegible] ```pythondef load_model(model_path: str): assert os.path.exists(model_path), f"Model path {model_path} invalid" return torch.load(model_path)```2. [illegible]```python# [illegible][bs, 2048]# [illegible][bs, 768]fusion_input = torch.cat([img_feat, text_feat], dim=1) # [illegible]``` [illegible]```python# [illegible] self.img_proj = nn.Linear(2048, 512)self.text_proj = nn.Linear(768, 512)fusion_input = torch.cat([self.img_proj(img_feat), self.text_proj(text_feat)], dim=1)``` [illegible] 1. [illegible]-[illegible] ```python# [illegible] def load_video(video_path): frames = [extract_frame(video, i*0.1) for i in range(100)] # [illegible] 0.1 [illegible] audio = extract_audio(video_path) # [illegible] return frames, audio``` [illegible]```pythondef load_video(video_path): cap = cv2.VideoCapture(video_path) fps = cap.get(cv2.CAP_PROP_FPS) frame_interval = 1/fps frames = [] timestamps = [] while cap.isOpened(): ret, frame = cap.read() if not ret: break frames.append(frame) timestamps.append(cap.get(cv2.CAP_PROP_POS_MSEC)/1000) audio = extract_audio_with_timestamps(video_path, timestamps) return frames, audio``` [illegible] 1. [illegible]```python# [illegible] GPU [illegible] loss = model(input)loss.backward()# [illegible] optimizer.step()``` [illegible]```pythonfrom torch.nn.parallel import DistributedDataParallel as DDPmodel = DDP(model)for device_input in distributed_inputs: loss = model(device_input) loss.backward() optimizer.step() optimizer.zero_grad()``` [illegible] 1. [illegible] ```pythondef process_data(): video = load_4k_video() # [illegible] 2GB [illegible] # [illegible]``

`[illegible]```pythondef process_data(): with VideoProcessor(video_path) as vp: frames = vp.get_frames() # [illegible] with [illegible] vp.release()```[illegible]"[illegible]"[illegible]1. [illegible]```python# [illegible] print("[illegible]:", video_feat.shape) # [illegible][batch, time, feat_dim]print("[illegible]:", question) print("[illegible]:", text_feat.mean().item())```2. [illegible]```python# [illegible]

●[illegible]- [illegible] [illegible] Python [illegible] `def myFunction():` [illegible] `def my_function():`[illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible] RGB [illegible] 224x224[illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible]RNN[illegible]LSTM[illegible]GRU[illegible]CNN[illegible]- [illegible] [illegible]- [illegible] [illegible] F1 [illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible] CPU[illegible]GPU [illegible]- [illegible] [illegible]

●[illegible]- [illegible] SonarQube[illegible]ESLint[illegible]Pylint [illegible] ESLint [illegible] JavaScript [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]

□□□- □□□□□ - □□ - □□- □□□□□ - GPU□□□□□□□□□□□□□□□□□□□□□□□□□ GPU□□ NVIDIA □ A100 □ H100 □□ - CPU□□□□□□□□□□□□□□□□□□□□ CPU□□ Intel Xeon Scalable □□□ AMD EPYC □□□□□□□□□□□□I/O □□□□□□□□□□□□□ - □□□□□□□□□ 128GB □□ 256GB □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ - □□□□□□□□□□□□□□□□□□ NVMe SSD□□□□□□ 1TB □□□□□□□□□□□□checkpoint□□□□□□□□□□□□□□□□□□□□□□

●□□---### □□□□□□□□□□□□□□□□□□□#### **1. □□□□□□□****□□ 1□N+1 □□□□**```python# □□□□□□□□□□ SQL □□ for user_id in active_users: profile = db.execute(f"SELECT * FROM user_profiles WHERE id={user_id}") # □□□□□□□□□□```**□□□□□**```python# □□□□□□ user_ids = [u.id for u in active_users]profiles = db.execute("SELECT * FROM user_profiles WHERE id IN %s", (tuple(user_ids),)profile_map = {p.id: p for p in profiles}```**□□ 2□□□□□□□□**```python# □□□□□□□□□□□□□□ update_account(balance-100)create_order(amount=100) # □□□□□□□□□□□□□□```**□□□□□**```pythonwith db.transaction(): # □□□□□□ update_account(balance-100) create_order(amount=100)```---#### **2. □□□□□□□□□□□****□□ 3□□□□□□□□□□□□□**```sql-- □□□□□□□□□□□□□ BLOBCREATE TABLE video_data (id INT PRIMARY KEY, raw_frames BLOB -- □□□□□□□□□);```**□□□□□**```sql-- □□□□□□ CREATE TABLE video_metadata (id INT PRIMARY KEY, s3_path VARCHAR(512), -- □□□□□□□□□ feature_vector VECTOR(2048) -- □□□□□□□□□);-- □□□□□□□□□□□ CREATE INDEX video_feature_idx ON video_metadata USING ivfflat (feature_vector);```---### □□□□□□□□□#### **1. □□□□□□**| □□□ | □□□□ | □□□□ ||------------------|-----------------------------|---------------------------------|| □□□□ | □□□□□□□□□□□□□ | □□□ users □ orders □ || □□□□ | WHERE create_time > ? □□□ | □□□□□□(create_time, status)|| □□□□□□ | □ VARCHAR □□ JSON □□ | □□ JSONB □□ |#### **2. □□□□□□**```sql-- □□□□□□ EXPLAIN ANALYZE SELECT * FROM multimodal_data WHERE created_at BETWEEN '2023-01-01' AND '2023-12-31';-- □□□ SELECT * FROM pg_locks WHERE relation = 'multimodal_data'::regclass;```#### **3. □□□□□□□□□□□**```pythondef validate_video_text_pair(video_id, text_id): # □□□□□□□ video_ts = db.execute("SELECT timestamps FROM video WHERE id=?", video_id)[0] text_ts = db.execute("SELECT timestamps FROM text WHERE id=?", text_id)[0] # □□□□□□□□□□□ dtw_distance = dtw(video_ts, text_ts) assert dtw_distance < 0.1, "□□□□□□□□" # □□□□□□□□□□ video_feat = db.execute("SELECT feature FROM video_features...") text_feat = db.execute("SELECT embedding FROM text_embeddings...") cos_sim = cosine_similarity(video_feat, text_feat) assert cos_sim > 0.7, "□□□□□□□"```---### □□□□□□□□#### **1. □□□□□□**| □□ | □□□□ | □□□□ ||----------------|---------------------------|---------------------------------|| **□□□□** | 2x RTX 4090 (24GB VRAM) | 8x A100 80GB + NVLink □□ || **□□** | 128GB DDR5 | 1TB DDR5 + 4TB Optane □□□□ || **□□** | 4TB NVMe SSD | 100TB All-Flash SAN + □□□□ || **□□** | 10Gbps □□□ | 100Gbps InfiniBand RDMA |#### **2. □□□□□□□ **```yaml# □□□□□□□□□□□□□□□cpu: 2x Intel Xeon Platinum 8480+ (56 □/□□□)memory: 2TB DDR5storage: - □□□: 30TB NVMe SSD (RAID 10) - □□□: 500TB QLC SSDgpu: 4x A30 (□□□□□□□□□)network: □ 100Gbps □□□□```#### **3. □□-□□□□□□**1. **□□□□□□□□□□**``` □□□□(G) = (□□□□□ × 4 □□) / 1e9 × □□□□(≥2.5) □□175B □□

[illegible] ≥ 175×4×2.5 / 1000 = 1.75TB [illegible] ```2. **[illegible]** ``` [illegible] IOPS = [illegible] × ([illegible] IOPS + [illegible] IOPS) [illegible]100 [illegible] × (3 [illegible]+1 [illegible]) = 400 IOPS [illegible]≥500 IOPS [illegible] ```---# ## [illegible]**[illegible]** [illegible]**[illegible]**1. **[illegible]** ```bash nvidia-smi # [illegible] GPU [illegible] iostat -x 1 # [illegible]>100ms ```2. **[illegible]** ```sql EXPLAIN (ANALYZE, BUFFERS) SELECT * FROM multimodal_index WHERE feature <-> ARRAY[...] < 0.2 ORDER BY similarity LIMIT 100; -- [illegible] ```3. **[illegible]** ```python # [illegible] py-spy [illegible] py-spy top --pid 12345 # [illegible] 70%[illegible] JSON [illegible] ```**[illegible]**1. [illegible]`CREATE INDEX CONCURRENTLY ON ... USING ivfflat`2. [illegible] JSON [illegible] Protocol Buffers [illegible] 3. [illegible] Redis [illegible]--- ### [illegible]```mermaidgraph TD A[[illegible]] --> B{[illegible]} B -->|[illegible]| C[[illegible]] B -->|[illegible]| F[[illegible]] C --> D[[illegible]] D --> E[[illegible]] E -->|[illegible]| G[[illegible]] E -->|[illegible]| H[[illegible]] subgraph [illegible] B1[[illegible]] B2[SQL [illegible]] B3[[illegible]] end subgraph [illegible] C1[[illegible]] C2[GPU [illegible]] C3[[illegible]] end```[illegible]**[illegible]**[illegible]

●[illegible]- [illegible] SonarQube[illegible]ESLint[illegible]Pylint [illegible] ESLint [illegible] JavaScript [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible] - [illegible] - [illegible]- [illegible] - GPU[illegible] GPU[illegible] NVIDIA [illegible] A100 [illegible] H100 [illegible] - CPU[illegible] CPU[illegible] Intel Xeon Scalable [illegible] AMD EPYC [illegible]I/O [illegible] - [illegible] 128GB [illegible] 256GB [illegible] - [illegible] NVMe SSD[illegible] 1TB [illegible]checkpoint[illegible]

●[illegible]- [illegible] [illegible] EXPLAIN[illegible] MySQL[illegible]PostgreSQL [illegible] MySQL [illegible] SELECT [illegible] EXPLAIN [illegible]- [illegible] SQL [illegible] SELECT [illegible] WHERE [illegible] WHERE [illegible] IN [illegible] OR[illegible]- [illegible] MySQL [illegible]- [illegible]- [illegible] TINYINT [illegible] INT[illegible] VARCHAR[illegible]- [illegible] B [illegible]- [illegible]

[illegible]- [illegible] [illegible]- [illegible] [illegible]- [illegible] [illegible] MySQL [illegible] innodb_buffer_pool_size [illegible] max_connections [illegible]- [illegible] [illegible] MySQL [illegible] PostgreSQL [illegible]- [illegible] [illegible]- [illegible] [illegible]

●[illegible] 1. [illegible] [illegible] 2. [illegible] [illegible] 3. [illegible] [illegible] 4. [illegible] [illegible] 5. [illegible] [illegible] MySQL 1. [illegible] [illegible] `EXPLAIN` [illegible]```sqlEXPLAIN SELECT * FROM table_name WHERE column_name = value;``` 2. [illegible] [illegible] - `id`[illegible] SELECT [illegible]id [illegible] - `select_type`[illegible] SIMPLE[illegible] UNION[illegible]SUBQUERY[illegible]UNION[illegible] - `table`[illegible] - `partitions`[illegible] - `type`[illegible] ALL[illegible]index[illegible]range[illegible]ref[illegible]eq_ref[illegible]const[illegible]system[illegible]type [illegible] - `possible_keys`[illegible] - `key`[illegible] - `key_len`[illegible] - `ref`[illegible] - `rows`[illegible] - `filtered`[illegible] 100 [illegible] - `Extra`[illegible] Using index[illegible]Using where[illegible] WHERE [illegible]Using temporary[illegible]Using filesort[illegible]PostgreSQL 1. [illegible] [illegible] `EXPLAIN` [illegible]```sqlEXPLAIN SELECT * FROM table_name WHERE column_name = value;``` 2. [illegible] [illegible] - `Seq Scan`[illegible] - `Index Scan`[illegible] - `Index Only Scan`[illegible] - `Bitmap Heap Scan`[illegible] - `Bitmap Index Scan`[illegible] - `Sort`[illegible] - `Aggregate`[illegible] - `Hash Join`[illegible] - `Merge Join`[illegible] - `Nested Loop`[illegible]SQL Server 1. [illegible] [illegible] SQL Server Management Studio (SSMS) [illegible] - [illegible] “[illegible]” [illegible] Ctrl + M [illegible] “[illegible]” [illegible] - [illegible] `SET SHOWPLAN_XML ON` [illegible] `SET SHOWPLAN_TEXT ON` [illegible] 2. [illegible] [illegible] - `Table Scan`[illegible] - `Clustered Index Scan`[illegible] - `Index Seek`[illegible] - `Bookmark Lookup`[illegible] - `Nested Loops`[illegible] - `Merge Join`[illegible] - `Hash Match`[illegible] - `Sort`[illegible] - `Filter`[illegible] - `Aggregate`[illegible] - `Compute Scalar`[illegible]

●[illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]

● ## [illegible] 1[illegible] ### [illegible] [illegible]+[illegible] ### [illegible] - **[illegible]**[illegible]Google Gemini [illegible] 23%[illegible]2025[illegible] - **[illegible]**[illegible]Meta [illegible] Llama 3 [illegible] Tensor [illegible] - **[illegible]**[illegible] CLIP[illegible]vs [illegible] ### [illegible] - **[illegible]**[illegible] 15%[illegible] ---## [illegible] 2[illegible] ### [illegible] OpenAI [illegible] Agent [illegible] ### [illegible] - **[illegible]**[illegible] GEA [illegible] 90%[illegible]2025[illegible] - **[illegible]**[illegible]70%[illegible] API [illegible] - **[illegible]**[illegible] vs [illegible] ### [illegible] - **[illegible]**[illegible] SQL [illegible] ---## [illegible] 3[illegible] ### [illegible] Faiss [illegible] ### [illegible] - **[illegible]**[illegible] PAI [illegible] LSH[illegible]<1%[illegible] - **[illegible]**[illegible] k-[illegible] - **[illegible]**[illegible] vs [illegible] ### [illegible] - **[illegible]**[illegible] 8%[illegible] ---## [illegible] 4[illegible] ### [illegible] GPU/TPU [illegible] ### [illegible] - **[illegible]**[illegible]NVIDIA H100 [illegible] FP8 [illegible] 4 [illegible] - **[illegible]**[illegible] PUE [illegible] 1.05[illegible] 30%[illegible] - **[illegible]**[illegible] H100 80GB[illegible] vs [illegible] ### [illegible] - **[illegible]**[illegible]+[illegible]BERT [illegible] 12 [illegible] 3.5 [illegible] 3B[illegible] ---## [illegible] 5[illegible] ### [illegible]A/B [illegible] ### [illegible] - **[illegible]**[illegible]DeepMind [illegible] AlphaTest [illegible] 60%[illegible] - **[illegible]**[illegible] AI [illegible] 20%[illegible] - **[illegible]**[illegible] vs [illegible] ### [illegible] - **[illegible]**[illegible]LangChain [illegible]Chain-of-Thought[illegible] 30%[illegible] ---## [illegible]5 [illegible] 1. **[illegible]**[illegible] Gemini[illegible] 2. **[illegible]**[illegible]70%[illegible] API [illegible] 3. **[illegible]**[illegible]+[illegible] 4. **[illegible]**[illegible]+[illegible] TCO 40%[illegible] 5. **[illegible]**[illegible]+[illegible] ---## [illegible] 1. **[illegible]**[illegible]Scaling Laws for Neural Language Models[illegible]OpenAI, 2025[illegible] 2. **[illegible]**[illegible]Hugging Face Transformers[illegible] 3. **[illegible]**[illegible]PyTorch Distributed Data Parallel[illegible]DDP[illegible] 4. **[illegible]**[illegible] GEA [illegible]GitHub[illegible] 5. **[illegible]**[illegible]LangChain Agent [illegible]

●[illegible] AI [illegible] json [illegible] 69%[illegible] 300 [illegible] Web [illegible] AKSK [illegible] SQL [illegible] OSS - Fuzz [illegible] OSS - Fuzz [illegible] 26 [illegible] OpenSSL [illegible]CVE - 2024 - 9143[illegible] Big Sleep [illegible] Big Sleep [illegible] Project Zero [illegible] DeepMind [illegible] SQLite [illegible] SQLite [illegible] prompt[illegible] SQLite [illegible]

●□□□---### □□□□□□□□□□□□□#### 1.1 □□□□□□□□□**□□□□**□□□□□□□□□□□□□□□□□```python# □□(16kHz)□□□(30fps)□□□□ audio = load_audio("clip.wav") # 16000 samples/secvideo = load_video("clip.mp4") # 30 frames/sec```**□□□□**□```pythondef align_multimodal(audio, video): # □□□□□□□□□□□□ video_timestamps = [i/30 for i in range(len(video))] # □□□□□□□□□□□□ audio_resampled = [] for ts in video_timestamps: start = int(ts * 16000) end = int((ts + 1/30) * 16000) audio_resampled.append(audio[start:end].mean()) return video, np.array(audio_resampled)```#### 1.2 □□□□□□□**□□□□**□□□□□□□□□□□□□□□□□```python# □□□□□□□□□□□□□ new_vision_model = VisionTransformer() # □□ 768 □ old_vision_model = ResNet50() # □□ 2048 □```**□□□□**□```pythonclass ModelRouter: def __init__(self): self.dim_map = { "vit": 768, "resnet": 2048 } self.proj_layers = nn.ModuleDict({ "vit": nn.Identity(), "resnet": nn.Linear(2048, 768) }) def get_features(self, model_type, x): feat = self.models[model_type](x) return self.proj_layers[model_type](feat)```---### □□□□□□□□□□#### 2.1 □□□□□□□□□**□□□□**□□□□□□□□□□□□```sql-- □□□□□□ 4K □□□□□□ BLOBCREATE TABLE media_data (id INT PRIMARY KEY, video BLOB, -- □□□□□□□□□ GB audio BLOB);```**□□□□**□```sql-- □□□□□□□□□□□□+□□□ CREATE TABLE media_metadata (id INT PRIMARY KEY, video_path VARCHAR(256), audio_path VARCHAR(256), duration FLOAT, resolution VARCHAR(16), created_at TIMESTAMP DEFAULT CURRENT_TIMESTAMP);-- □□□□□□□ CREATE TABLE media_features (id INT REFERENCES media_metadata(id), frame_features JSONB, # □□□□□□ audio_features FLOAT[]);```#### 2.2 □□□□□□**□□□□**□□□□□□□□□□□□```sql-- □□□□□□□□ SELECT * FROM images JOIN text_data ON images.timestamp = text_data.timestampWHERE text_data.keyword = 'emergency';```**□□□□**□```sql-- □□□□□□ CREATE INDEX idx_multimodal_search ON text_data USING gin(keyword gin_trgm_ops, timestamp);-- □□□□□ CREATE TABLE text_data_2023 PARTITION OF text_dataFOR VALUES FROM ('2023-01-01') TO ('2024-01-01');```---### □□□□□□□□□#### 3.1 □□□□□□□□□| □□ | □□□□ | □□□□□□□ ||------|----------|----------------|| GPU | NVIDIA A100 80GB x8 | NVIDIA H100 80GB x32 || CPU | AMD EPYC 7763 64 □ | 2x Intel Xeon Platinum 8480+ || □□ | 512GB DDR5 | 2TB DDR5 ECC || □□ | 4x NVMe SSD 3.84TB RAID0 | Ceph □□□□□ 1PB || □□ | 100GbE | 400Gb InfiniBand |#### 3.2 □□□□□□□□□| □□ | □□ ||------|------|| □□□□ | PostgreSQL 14 + TimescaleDB (□□□□) || □□ | □ TB □□□□ 128GB □□ || □□ | NVMe SSD RAID10□□□ 30% OP □□ || □□□ | □□□□□ x3 □□(max_connections) || □□ | □□ WAL □□ + □□□□□□□□□□□□ |---### □□□□□□□□□□**□□□□**□□□□□□□□□□□□□□□□□#### 4.1 □□□□□□□□□1. **□□□□□□**```sqlEXPLAIN ANALYZE SELECT video_path FROM media_data WHERE similarity_rank(audio_features, '[0.12, 0.34,...]') > 0.8ORDER BY timestamp DESC LIMIT 10;```□□□□□□□□ 2.3 □2. **□□□□**□```sql-- □□□□□□ CREATE INDEX idx_audio_features ON media_data USING ivfflat (audio_features vector_cosine_ops) WITH (lists = 1000);-- □□□□□□ 87ms```#### 4.2 □□□□□□□□□□```bash# □□ GPU □□□□ nvidia-smi --query-gpu=memory.used --format=csv -l 1# □□□□□□□□□□□□□# □□□□□□□□□□□□```---### □□□□□□□□□□□#### 5.1 □□□□□□□□□□```python# □□□□□□□ import psycopg2conn = psycopg2.connect(dbname="multimodal_db")cur = conn.cursor()cur.execute("""SELECT schemaname, tablename, indexname, round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid),2) as

index_ratioFROM pg_indexes WHERE schemaname = 'public'ORDER BY pg_relation_size(indexrelid) DESC;""")for row in cur.fetchall(): if row[3] > 30: # □□□□□□□□□ 30% print(f"□□: □□{row[2]}□□□□□□□")```#### 5.2 □□□□□□□□```yaml# Prometheus □□□□□ scrape_configs: - job_name: 'gpu_metrics' static_configs: - targets: ['gpu-exporter:9100'] - job_name: 'db_metrics' params: collect: - 'custom_queries' static_configs: - targets: ['postgres-exporter:9187']```---### □□□□□□ checklist1. **□□□□□**□ - □□□□□□□□□□□□□□/□□□□□ - □□□□□□□□□□□IVFFlat vs HNSW□ - □□□□□□□□□□□= (CPU □□□ * 2) + □□□□□□2. **□□□□**□ - GPU □□□□□□□□□□□ (GB) x 4.5□□□□□+□□□□□□ - □□□□□□□□□□□□□ x 1.5 □□ - □□□□□□□□□□□□□□□□□□□ < □□□□□ 20%3. **□□□□**□ ```python # □□□□□□□□□ def safe_db_write(conn, query, max_retries=3): for _ in range(max_retries): try: conn.execute(query) conn.commit() return True except psycopg2.OperationalError as e: conn.rollback() reset_connection(conn) return False ```□□□ 30%□□ 72 □□□□□□□□□□□□□□□□□□ GPU □□---### □□□□□□□□□□□□□#### 1.1 □□□□□□□□□**□□□□**□□□□□□□□□□□□□□□□□```python# □□(16kHz)□□□(30fps)□□□□ audio = load_audio("clip.wav") # 16000 samples/secvideo = load_video("clip.mp4") # 30 frames/sec```**□□□□**□ ```pythondef align_multimodal(audio, video): # □□□□□□□□□□□□ video_timestamps = [i/30 for i in range(len(video))] # □□□□□□□□□□□□ audio_resampled = [] for ts in video_timestamps: start = int(ts * 16000) end = int((ts + 1/30) * 16000) audio_resampled.append(audio[start:end].mean()) return video, np.array(audio_resampled)```#### 1.2 □□□□□□□**□□□□**□□□□□□□□□□□□□□□□ ```python# □□□□□□□□□□□□□ new_vision_model = VisionTransformer() # □□ 768 □ old_vision_model = ResNet50() # □□ 2048 □```**□□□□**□```pythonclass ModelRouter: def __init__(self): self.dim_map = { "vit": 768, "resnet": 2048 } self.proj_layers = nn.ModuleDict({ "vit": nn.Identity(), "resnet": nn.Linear(2048, 768) }) def get_features(self, model_type, x): feat = self.models[model_type](x) return self.proj_layers[model_type](feat)```---### □□□□□□□□□#### 2.1 □□□□□□□□□**□□□□**□□□□□□□□□□□```sql-- □□□□□ 4K □□□□□ BLOBCREATE TABLE media_data (id INT PRIMARY KEY, video BLOB, -- □□□□□□□□ GB audio BLOB);```**□□□□**□```sql-- □□□□□□□□□□□+□□□ CREATE TABLE media_metadata (id INT PRIMARY KEY, video_path VARCHAR(256), audio_path VARCHAR(256), duration FLOAT, resolution VARCHAR(16), created_at TIMESTAMP DEFAULT CURRENT_TIMESTAMP);-- □□□□□□ CREATE TABLE media_features (id INT REFERENCES media_metadata(id), frame_features JSONB, # □□□□□ audio_features FLOAT[]);```#### 2.2 □□□□□□**□□□□**□□□□□□□□□□□```sql-- □□□□□□□□□ SELECT * FROM images JOIN text_data ON images.timestamp = text_data.timestampWHERE text_data.keyword = 'emergency';```**□□□□**□ ```sql-- □□□□□ CREATE INDEX idx_multimodal_search ON text_data USING gin(keyword gin_trgm_ops, timestamp);-- □□□□ CREATE TABLE text_data_2023 PARTITION OF text_dataFOR VALUES FROM ('2023-01-01') TO ('2024-01-01');```---### □□□□□□□□#### 3.1 □□□□□□□□| □□ | □□□□ | □□□□□□ ||------|----------|----------------|| GPU | NVIDIA A100 80GB x8 | NVIDIA H100 80GB x32 || CPU | AMD EPYC 7763 64 □ | 2x Intel Xeon Platinum 8480+ || □□ | 512GB DDR5 | 2TB DDR5 ECC || □□ | 4x NVMe SSD 3.84TB RAID0 | Ceph □□□□ 1PB || □□ | 100GbE | 400Gb InfiniBand |#### 3.2 □□□□□□□□| □□ | □□ ||------|------|| □□□□ | PostgreSQL 14 +

TimescaleDB ([illegible]) || [illegible] | [illegible]TB [illegible] 128GB [illegible] || [illegible] | NVMe SSD RAID10[illegible] 30% OP [illegible] || [illegible] | [illegible] x3 [illegible](max_connections) || [illegible] | [illegible] WAL [illegible] + [illegible] |--- ### [illegible]**[illegible]**[illegible]#### 4.1 [illegible]1. **[illegible]**```sqlEXPLAIN ANALYZE SELECT video_path FROM media_data WHERE similarity_rank(audio_features, '[0.12, 0.34,...]') > 0.8ORDER BY timestamp DESC LIMIT 10;```[illegible] 2.3 [illegible] 2. **[illegible]**[illegible]```sql-- [illegible] CREATE INDEX idx_audio_features ON media_data USING ivfflat (audio_features vector_cosine_ops) WITH (lists = 1000);-- [illegible] 87ms```#### 4.2 [illegible]```bash# [illegible] GPU [illegible] nvidia-smi --query-gpu=memory.used --format=csv -l 1# [illegible]# [illegible]```---### [illegible]#### 5.1 [illegible]```python# [illegible] import psycopg2conn = psycopg2.connect(dbname="multimodal_db")cur = conn.cursor()cur.execute("""SELECT schemaname, tablename, indexname, round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid),2) as index_ratioFROM pg_indexes WHERE schemaname = 'public'ORDER BY pg_relation_size(indexrelid) DESC;""")for row in cur.fetchall(): if row[3] > 30: # [illegible] 30% print(f"[illegible]: [illegible]{row[2]}[illegible]")```#### 5.2 [illegible]```yaml# Prometheus [illegible] scrape_configs: - job_name: 'gpu_metrics' static_configs: - targets: ['gpu-exporter:9100'] - job_name: 'db_metrics' params: collect: - 'custom_queries' static_configs: - targets: ['postgres-exporter:9187']```---### [illegible] checklist1. **[illegible]**[illegible] - [illegible]/[illegible] - [illegible]IVFFlat vs HNSW[illegible] - [illegible]= (CPU [illegible] * 2) + [illegible]2. **[illegible]**[illegible] - GPU [illegible] (GB) x 4.5[illegible]+[illegible] - [illegible] x 1.5 [illegible] - [illegible] < [illegible] 20%3. **[illegible]**[illegible] ```python # [illegible] def safe_db_write(conn, query, max_retries=3): for _ in range(max_retries): try: conn.execute(query) conn.commit() return True except psycopg2.OperationalError as e: conn.rollback() reset_connection(conn) return False ```[illegible] 30%[illegible] 72 [illegible] GPU [illegible]

●[illegible]- [illegible] L4 [illegible] 28%[illegible] 2000 [illegible]- [illegible] L4 [illegible] 360 [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible]- [illegible] GPT-4 [illegible]- [illegible]- [illegible]- [illegible]- [illegible] - [illegible] - [illegible] - [illegible] - [illegible]

- □□□ □□□□ - □□□□ □□□□□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□□□□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□ - □□□□ □□□□□1 - 10 □□ 1 - 20 □□ 1 - 10 □□ □□

●□□L4 □□□□□□□- □□□□□ □L4 □□□□□□□□□□□□□□□ 6-12 □□□□□3-12 □□□□□□□5 □□□□□□□□□□□ 1-2 □ GNSS/IMU □ 1-2 □□□□□□□□□□□□□□□□□□□□□□□□□□□□ 10 □□□□2025 □□□□ 800-1500 □□ / □□□□□□□ 30%-40%□□□□□□□□□□□□□□□□□□□ 20%-30%□□□□□□□□□□ 15-20 □□ / □□□□□□□□□□- □□□□□□ □□□□□□□□□ Orin □□□□□□□□□□□ 8000 □□□□□□□□ 25%-35%□- □□□□□□ □□□□□□□□ 2000-5000 □□ / □□□□□□□ 10%-15%□- □□□□□□□ □□□□□□L4 □□□□□□□□□□□□□□□□□ 50% □□□□□□□□ L4 □□□□□□□□□□□□□□□□□□□□□□□□□ 35%□L2 □□□□□□□- □□□□□ □L2 □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ 1000 □□□□□□□□□77GHz□□ 1000 □□ / □□□□□□□□□ 20%-30%□- □□□□ □□□□□□□□□□□ L2 □□□□□□□□□□□□□□□ 25%-35%□□□□□2025 □□□□□□□□□□□□□□□□□□2026 □□□□□ 50%□- □□□□□□□ □L2 □□□□□□□□□□□□□□□□□□□□□□□□□□□□ 35%-45%□

●□□□---### **□□□□□□□□□□□□□□□□□□□**#### 1. **□□□□□□□□□□□**- **□□**□8000 □-2 □□ - **□□□□□□□□**□2000 □-5000 □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ L4 □□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□**□4000 □-1.2 □□□□□□□□□□□□□□□□□□□□□□3.8 □/□□□□□□□+□□ GPU □□□□□□□□□□□□□□□□□ - **□□□□**□2000 □-3000 □□□□□□□□□□□□□□□□□ FSD □□□□□□□□□□□□□□□□□□□□□ #### 2. **□□□□□□□□□□□□**- **□□**□3000 □-8000 □□ - **□□□□□□□**□□□□-□□□□□□□□□□□□□□□□□□□□□□□ V2X 5.0 □□□ 500 □□□□□□□□□□□□ - **□□□□□□**□□□ NVIDIA Omniverse □□□□□□□□□□□□□□□□□□□□□□□□ 70%□□□□□□□□□□□□□□□ 50% ---### **□□□□□□□□□□□**#### **1. □□□□□□**- **□□□□**□ - **□□□□□**□20 □-30 □□/□□□□□□□□□□ 3 □□/□□□ 128 □□□□□□□□□□□ - **□□□□**□5 □-8 □□/□□□□□□□ 6 □□+□□□□□□□□□□□□1000TOPS□ - **□□□□□**□ - **□□□□□□□□**□1000 □-2000 □□□□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□**□800 □-1500 □□□ASIL-D □□□□□□+□□□□□□□□□ - **□□□**□□□□□□ 2.5 □-6 □□□□100 □□□□□+□□□□□ #### **2. □□□□□□□□**- **□□□□□□□□□**□ - **□□□□**□4 □-6 □□□□□□□□□□□□□+□□□□□□□□ 35%-50%□ - **□□□□**□1 □-2 □□□4D □□□□□+IMU□□□□□□□□ 68%□□□ - **□□□**□3.5 □-5 □□□□□□□□□+□□□□□□□□□□□□□□□□□□□□ 30%□ - **□□□□□**□ - **□□□□□□**□5000 □-1 □□□□□ NVIDIA GR00T-Dreams □ 36 □□□□□□□□□□□□□□ 3 □□□□□ - **□□□□□□□**□□□ 25 □-35 □□□□□ 2027 □□□ 15 □□□□□ - **□□□**□3 □-8 □□□□□□□□□+□□□□□□ ####**3. □□□□□□□□□□**- **□□□□□□□□□**□ - **□□□□**□8000-1.5 □□□□□□□□□□+□□□□□□□□□□□ - **□□□□**□2 □-3 □□□SLAM □□+□□□□□□□□□□□□□□ 40%□ - **□□□□□**□ - **□□□□□**□800 □-1500 □□□□□□□□□□□□□□□□□ - **□□□□□**□500 □-1000 □□□GLM-PC 1.1 □"□□□□+□□□□"□□□□□□ - **□□□**□1.2 □-3 □□□□□□□□□□+□□□□□□ ---### **□□□□□□□□□□□□**#### 1. **□□□□□**- **□□□□**□1200 □-2500 □□□50 □□□□□□□□□□□+□□□ V2X □□□ - **□□□□□**□600 □-1000 □□□CE/FCC □□+□□□□□□□□□□ #### 2. **□□□□**- **B □□□**□2000 □-4000 □□□□□□/□□□□□□□□□□"□□□□+□□□□"□□□ - **C □□□**□1500 □-3000 □□□□□/□□□□□□+□□□□□□□□ Gokit5 □□□□□□□□□ #### 3. **□□□□**- **OTA □□**□□□□□□□□ 8%-12%□□□□□□□10TB □□□□□□□ - **□□□□**□□□□□□□□ 10%-15%□□□□□□□ 60 □□□□□□ ---### **□□□□□□□□□□□□**1. **□□□□□□**□ - □□□□/□□□□□□□□□□□□□□□□□□□□□□□□ 5%□□ 2. **□□□□□□**□ - □□□□□"□□□□□□□□□□□□□□"□□□ 30%-50%□□□□ 3. **□□□□**□ - □□□□□□□□□□□□□□□□□□□□□ 4D □□□□□□□□□□□ ---### **□□□□□□□□□□□□□□□**| **□□** | 2025 □□□□□ | 2027 □□□□□ | □□□□□2030 □□ | |----------------|----------------|----------------|--------------------| | □□□□□□ | 20 □□ | 12 □□ | 120 □□□□□□ | | □□□□□ | 35 □□ | 15 □□ | 5900 □□□□□□ | | □□□□□ | 3 □□ | 1.2 □□ |

3000 □□□□□ | **□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□---### **□□□□□**- **□□□□**□□□□□□□□□□□□□20 □□□□□□□70% - **□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□50%□□35% - **□□□□□**□Gokit5 □□□□□“□□□□+□□□□□”□□□□□□□□10 □□□ □□□

●□□□---### **□□□□□□□□□□□□□□□□□**1. **□□□□□□□□** - **□□**□□5000 □-1.5 □□ - **□□**□□GPU/TPU □□□ - **□□□**□ - □□□□□□□□□1000 □-3000 □□ - □□□□□□GPU □□□3000 □-1 □□ - □□□□□□□□□1000 □-2000 □□ 2. **□□□□□□□□□□□** - **□□**□2000 □-5000 □□ - **□□**□□□□□□□□□□□-□□3. **□□□□□□** - **□□**□1 □-3 □□□□□□□□□ - **□□**□□□AI □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□16vCPU □□□□3.8 □□□□□□□□□ ---### **□□□□□□□□□□□**#### **1. □□□□□□** - **□□□□**□ - **□□□□□**□□□□□8 □-20 □□/□□□□□□5000-2 □□/□□□□□□□□1 □-5 □□/□□□□□□□□□20 □-50 □□□ - **□□□□**□□AI □□□□□□Orin□5000-2 □□/□□+ □□□□□□□□□5 □-15 □□□ - **□□□□□□□**□ - **□□□□□□□□□**□1000 □-3000 □□ - **□□□□□□**□500 □-1500 □□ - **□□□□**□□ISO 26262□□500 □-1000 □□ - **□□□**□□□□□□2 □-5 □□□□□□□□□□□□□□□□ #### **2. □□□□□□** - **□□□□□**□□□□□2 □□□□□□□ - **□□□**□□□□□□14%□□□□□□□□□□9%□□□□□□19%□□□□□□□□□1.2 □-1.8 □□□□ - **□□□**□□□□□□□□3%□□□□□□3D □□+□□□□□□□□□2000-5000 □□□ - **□□□□**□FSD □□□39%□□□□□□□□□□□□□□□□□□□5000 □-1 □□□ - **□□□□**□ - **□□□□□□□□□**□□□□□3 □-5 □□□□□□□3000 □-5000 □□□□ - **□□□□□**□□□□□□□□□□□□□□□20%-30%□□□ #### **3. □□□□□□□□□** - **□□□□**□ - **□□□□**□□□□□500-2000 □/□□□□□□□□□1000-3000 □/□□□ - **□□□□**□SLAM □□□□□1 □-3 □□/□□□ - **□□□□□□□**□ - **□□□□□□□□□**□500 □-1000 □□ - **□□□□□□**□300 □-800 □□ - **□□□□**□□□□□2 □-5 □□□□□□1 □□□□□□□□□□□□ ---### **□□□□□□□□□□□**1. **□□□□□□□** - **□□□□**□□□□□□□□□□□NCAP□□1000 □-3000 □□□ - **□□□□□**□CE/FCC □□□□□□□□500 □-1000 □□□ 2. **□□□□□□□□□** - **□□□□**□2000 □-5000 □□□□□□□Optimus □□□□□□□□ - **B □□□**□□□□□□□□□□□□500 □-2000 □□□ 3. **□□□□□** - **OTA □□**□□□□500 □-1000 □□ - **□□□□**□□□□□□□□□10%-20%□ ---### **□□□□□□□□□□**| **□□** | **□□□□** | **□□□□□** | **□□□□□** | |------------------------|-------------------|------------------- |-------------------| | □□□□□□□□ | 40% | 35% | 25% | | □□□□□□□□ | 50% | 60% | 40% | | □□□□□□□ | 10% | 5% | 35% | | **□□□□□□** | 5 □-10 □□ | 3 □-8 □□ | 1 □-3 □□ | ---### **□□□□□□□□□□**1. **□□□□□□**□□□2. **□□□□□□**□□□□□□CPU+GPU □□□□□□□□□□□□□□□□□□16vCPU □□3.8 □/□□□□3. **□□□□**□□□□□AI □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ ---### **□□□□□□□□□**- **□□□□**□□□□FSD □□□□□□□□□□□□□□2 □□□/□□ - **□□□□□**□□□□Walker X □□□□□□□□□□□□30%□ - **□□□□□**□□□□□□□□+AI Agent □□□□□□□□□□□□□□□ □□marketsandmarkets □□□AI Agent □□□□□□43%□□□□□□□□□□

●●□□---### **□□□□□□□□□□□□□□**#### **1. □□□□□□□□□□□**- **□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□Emu3 □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□- **□□**□□OpenAI □Operator □□□□□□□-□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□#### **2. □□□□□□□□□**- **□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□PaLM-E □□□□□□□□□□□□□□□□□□□□□□□-

□□□□“□□□□”□□□□□□□□□□□□□□□□□□□□□#### **3. □□□□□□**- **□□**□□- **□□**□□#### **4. □□□□□□**- **□□**□AI □□□□□□□□□□□□□□□□□□□□□□□- **□□**□□□□□□□“□□□□”□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□---### **□□□□□□□□□**#### **1. □□□□□□**- **□□□□□“□□□□”**□□□□□□□□□□□□□BiCR-SLAM □□□□□□□□□□□□□□□□□□□□□□□□□□□□□AI □□□□□□□□□□- **□□□□□□□□□□**□□□□□□□□□□□□□□□□□□□□□□FSD □□□□□□□□□□- **□□□□□□□□□**□□□□□□“□□□□+□□□□”□□□□□□□□□□□□□□□□#### **2. □□□□□□**- **□□□□**□AI4S□□□□□□□□□□□□□□□□□□□□□AI □□□□□□□□□□□□□□□□□- **□□□□**□□□□□□□□□□□□□□□□+□□□□□□□□□“□□ 2.5”□□□□□□□□□- **□□□□□□□**□□□□□□□□□□□□□□□Walker X□□□□□□□□□□□□□□□□□GLM-PC □□□□□□□---### **□□□□□□□□□**#### **1. □□□□□□□□□**- **□□□□□□□**□□□□□□□□□□□□□□□□□□□□□□50%□□□□□□□□□□□□□□ 30%□- **□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□6□□□□□□□□#### **2. □□□□□□□□□**- **□□□□**□□□□□□□□3 □□/□□□□□□□□□NVIDIA Omniverse□□ 70%□□□□□□□L4 □□□- **□□□□□□**□□□ RAG□□□□□□□□□□□□□□□□□□□□□□□ 7000 □□□□□□□□□#### **3. □□□□□□**- **□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□- **□□□□**□□□□ AI □□□□□□□□□□□□□□□□□□□□□□□□---### **□□□□□□□□**1. **□□□□**□ - **□□**□DeepMind □ AlphaFold3 □□□□□□□□□□□□□□□□□□60%□ - **□□**□AI4S □□□□□□□□□□□□□□□□□□□□□□□2. **□□□□**□ - **□□**□NASA □□AI □□□□□□□□□□□□□□□□□5 □□ - **□□**□□□□□□□□□□□□□BiCR□□□□□□□□□□3. **□□□□□□**□ - **□□**□IBM Watson Oncology □□□□□□□□□□□□□□□□□□90%□ - **□□**□□□□□□□□□□□□□□□---### **□□□□□□□**□□5 □□□□□□□□□□“□□”□□□**□□□□**□□□□□□□□□**□□□**□□□□□□□□□□□**□□□**□□□□□□□□□□□□□□□□□□□□□**□□□□**□□□□□□□□□**□□□□**□□□□□+□□□□□□**□□□□**□□□□□□□□□2030 □□□□□□□□□□□□□676 □□□□AI □□□□□□□□□□□□□□□□

●### □□□□□□□□□□□□□□□□ □□ CEO/□□□□□□□□□□□□□□□□□□□ ---#### **1. □□□□□□□□□□□□□□** **□□**□ □□□□□ GPT-4□Manus□□ **□□□□**□ - **□□□□**□□□□□□□□□□□□□□□□□□□□□□□□ - **□□□□□**□□□ Gemini□OpenAI □□□□□□□□□□□□□□□□□□□□□□□/□/□/□□□□□□□□□□□□□□□□□□□ 3%□□ - **□□□□□□**□□□□□□□□AI □□□□□□□□□□90%□□□ 1 □□□ 284 □ 720P □□□ **□□□**□ - **□□□□□□□□□**□1750 □□□□□□□□□□□□□□□□□□□□□□AI □□□□ 30%□□ - **□□□□□□**□□□□□□□□□□□□2026 □□□MIT □□□□□□□□RAG □□□□□ ---#### **2. □□□□□□□□□□□□□□** **□□**□ □□□□□□ Manus□□□□□□□□□□□□□□□□□□□□□“□□□”□□“□□□□”□ **□□□□□**□ - **□□□□□**□□□□□□ 82%□□□ 3 □□□□□□□□□□□/□□□□□70%□□ - **□□□□**□AI □□□□□□□□□AlphaFold3 □□ 6 □+□□□□□□□□□□1 □□□ - **□□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□ **□□**□ - **□□□□□**□□□/□□□□□□ERP/EMR □□□□□□□□□□□□□□□ - **□□□□□**□□□□□□□□□□□□□□□□L4 □□□□□□□□15%□□ ---#### **3. □□□□□□□□□□□□□□** **□□**□ □□□□□□□□□□□□□□□□□□□□□□**□□□□**□ - **□□□□**□□□□□□□□□□□□□2024 □□□ AI □□□□□□□ 240%□□ - **□□□□**□□□/□□□□□□□□□□□□□□□□AI □□□□□□□□□ - **□□□□□**□“□□□□ AI”□□□□□□□□□ **□□□□**□ - □□□□+□□□□□□□□□□□Azure □□□□□□ - □□□□□□□□□□□□□□□□□□□□□□□ ---#### **□□□□** 1. **□AI □□□□□□□□□□□□**□CSDN □□□□ AutoGen/CrewAI □□□□□ [□□](https://blog.csdn.net/2401_85373691/article/details/145255211) 2. **□2025 □ AI □□□□□□□□**□□□□□24 □□□ 989 □□□□□□□□ [□□](http://sc.stock.cnfol.com/gushiyaowen/20250102/31091581.shtml) 3. **□□□□□

□“□□”□□□□□□**□□□□□□□□□+□□□□□□□ [□□](http://mrdx.xinhuanet.com/20250331/f3a62c231f8b4ddda66bb282de9f4341/c.html) 4. **□AI □□□□□□□**□□□□□□□□□□□□□□□□□□□□□ ---### **□□□□**□5 □□□□□□ 1. **□□□□**□2025 □□□□□□□□□□□□□□+□□□□□□□□□□□□□ 2. **□□□□**□□□/□□/□□□□□□□□□□□□□□□□□ 50%□□ 3 □□□□□□ 3. **□□□□**□□□□□□□□P800□□□□□□□□□□□□□□□□□□□□ 4. **□□□□**□□□□□□□AutoGen□□□□□□□□□□□□□□□□□□□□□□□□ AI □□□□ 5. **□□□□**□□□ AI □□ 2026 □□□□□□□□□□□□□□□□□4%□□□ --- **□□□□**□□□□□□□/□□□□□□□□□□□□□□□□□□□□□□□□ AI □□□□□□□□□□□□

●### □□□□□□□□□□□□□□□2025□---#### **1. □□□□□□□□□□□□□□□□****□□**□□□ AI □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ **□□□□**□ - **□□**□□□ AI □□□□□□ 43%□WIPO 2024□□□□□□□□□□□□□□□□□□□□ SkyReels-V1 □ SOTA □□□□□□□□□□□□□ AI+□□□□□□ 80%□□□□□ - **□□**□□□ 77%□□ AI □□□□□RAND 2025□□OpenAI □ GPT-5 □□□□ 3 □□□□□ Copilot □□□□ 2.6 □□□□□ - **□□**□□□□□□□□□□□□□□□□□□□□ AI □□□□□□□□□ AI □□□□□□ 25TOPS/W □□□□ 3nm □□□ **□□**□ - □□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ 0.3%□ - □□□□DeepSeek-R1□vs □□□□GPT-5□□□□□□□□ ---#### **2. □□□□□□□□□□****□□**□□□ AI □□□□□"□□□□"□□□□□□□□□□□□□□□ **□□□□**□ - □□ Top100 AI □□□ 58%□□□□□□□□□□□□□□□□□□□ 2018 □ 12%□□ 2025 □ 29% - □□ AI Lab □□□□□□□$1.2M□□□□□ DeepMind $0.8M□ - □□"□□□□□"□□□□ AI □□□□□□□ AI □□□□□□□ 47% **□□□□**□ - □□□□□ MBZUAI □□□□□□+□□□□□□□□□□ - □□"□□□□□□"□□ MIT □□□□□□$500K **□□**□ - □□□□□□□□□□□□□□□□□□□ - □□□□□□□□□□□□□□□□□ ---#### **3. □□□□□□□□□□□****□□**□AI □□□□□□"□□□□□"□□□□□□/□□□/□□□□ **□□□□**□ - □□ AI □□□□ 2025 □□$1.8 □□□IDC□□□□□ - □□□□□□/□□□□□□ 63%□□□□□□ H100 □□□ 78%□ - □□□□□□□□□□□□□□□□□ 2.3 □□HuggingFace □□□ - □□□□ - AI □□□□□□□ DramaWave □□□$1000 □□ - AI □□□SkyMusic □□ ARR $1200 □□ - □□□□□□□□□□□□□□ 40%□ **□□□□□□**□ - □□□"□□□□□"□MaaS□□□□□ 7 □□□□□□□ - □□□ Optimus □□□□□□□□□ 50 □□ **□□**□ - □□□□□□□□□□□□Meta Llama3 □□□□$3 □□ - □□□□□□□□□□□□□□□□□□□□□□□ 300%□ ---#### **□□□□□□**1. □□□□□□□□□□□□(2024)□□□□□□□□□□□□ 2. □□□□2025 □□□□□□□□□□□□□□□□□□□□□ 3. RAND□□□ AI □□□□□□□□□□□□□□ 4. □□□□ 2024 □□□AIGC □□□□□□□ 5. MacroPolo □□ AI □□□□□□□□□ ---### **□□□□□□□□□□□**1. **□□□□**□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□ 2. **□□□□**□□□□□□□ AI □□□□ 5 □□□□□□□□□ 58%□□□□ 3. **□□□□**□AI □□/□□□□□□□□□□□□□□□□□□$1000 □□□ 4. **□□□□**□□□□□□□□□□□ 0.3%□2026 □□□□□□□□ 5. **□□□□**□□□ 77%□□□□□□□□□□□□□□□□□ 12%□□□ 35% □□□□□ 2025 □ 5 □□□□□□□□□□□□□□□□□□□□Ai artificial intelligence big model multi-model multi-modal general agent research and development actual combat □ 2025 v1.4e-book. ● Artificial intelligence big model multi-model multi-modal universal intelligent body. The necessary knowledge reserve and skill requirements for scientists and senior technical experts: knowledge reserve-foundation of mathematics and statistics-linear algebra: used to handle vectorization and matrix representation of data, as well as various operations in the process of model training and optimization, such as matrix multiplication and eigenvalue decomposition. -Probability theory and mathematical statistics: provide theoretical basis for model uncertainty modeling, parameter estimation and hypothesis testing, and help to understand the distribution and laws of data. -Optimization theory: in model training, the loss function is minimized by optimization algorithm (such as gradient descent method) to improve the performance of the model. -Basic theory of artificial intelligence-Machine learning: master the basic concepts, algorithms and principles of supervised learning, unsupervised learning and

reinforcement learning, such as linear regression, decision tree, support vector machine, clustering algorithm and dimension reduction algorithm. -Deep learning: Understand the basic structure and working principle of neural network, including multilayer perceptron, convolutional neural network, cyclic neural network, Transformer architecture, and their applications in different tasks. -Reinforcement learning: It studies how agents learn the best strategies through interaction with the environment, including Q-learning, strategy gradient method, actor-critical algorithm, etc. -Multi-modal data processing and knowledge fusion-Characteristics and processing methods of multi-modal data: Understand the characteristics and representations of image, text, voice, video and other multi-modal data, and master the methods of preprocessing, feature extraction, feature alignment and other operations on these data. -Multi-modal fusion technology: Learn how to effectively fuse data of different modes to achieve more comprehensive and accurate information understanding and task decision-making, such as early fusion, late fusion, intermediate fusion and other strategies. -Knowledge about the big model-Architecture and principle of the big model: deeply understand the architecture design and working principle of the big language model and multi-modal big model, such as the self-attention mechanism in the Transformer architecture and the training objectives of the pre-training language model. -Pre-training and fine-tuning technology: master the pre-training methods and fine-tuning strategies of large models, including supervision fine-tuning, reinforcement learning fine-tuning, and how to fine-tune and optimize large models according to different task requirements. -Model optimization and compression technology: learn optimization and compression methods such as model pruning, quantification and distillation to improve the operation efficiency and adaptability of large models. -Domain knowledge-Specialized knowledge in a specific field: For the application fields of general agents, such as autonomous driving, humanoid robots, home advanced intelligent robots, etc., understand the professional knowledge and technical requirements in related fields. -Industry trends and trends: pay attention to the latest technical trends and development trends in the field of artificial intelligence, as well as changes in policies, regulations and market demand of related industries. Skills requirements-programming and software development ability-proficient in programming languages: proficient in Python, C++ and other programming languages, able to write and debug codes efficiently and realize the development of algorithms and models. -Familiar with deep learning frameworks: Skillfully use deep learning frameworks such as PyTorch and TensorFlow, as well as related tools and libraries, such as Hugging Face Transformers, DeepSpeed, Megatron-LM, etc., to quickly build and train models. -Code management and collaboration ability: master version control tools (such as Git), and be able to manage, version control and collaborate on code development to ensure the smooth progress of the project. -Large-scale model training and optimization ability-Model training and optimization: have the experience and ability of large-scale model training, and can train and optimize the model according to the task requirements, including parameter adjustment, optimization algorithm selection, loss function design, etc. -Distributed training and parallel computing: master distributed training technology and parallel

computing framework, such as Horovod and NCCL, and be able to use multi-GPU and multi-node for efficient large-scale model training. -Model compression and deployment: large models can be compressed and optimized to adapt to different hardware platforms and application scenarios, and model deployment and reasoning optimization can be carried out. -Data processing and analysis capabilities-Data collection and preprocessing: able to collect, clean, label and enhance large-scale multimodal data to ensure data quality and availability. -feature engineering and data analysis: feature extraction, feature selection and feature engineering are carried out on the data to improve the performance and generalization ability of the model. -Algorithm research and innovation ability-Algorithm design and improvement: Have the ability to design and improve artificial intelligence algorithms, and be able to propose effective solutions and algorithm innovation for practical problems. -Reading and reproducing papers: able to read and understand relevant papers in international top conferences and journals, quickly reproduce and verify new algorithms and models. -System architecture and engineering capability-System design and architecture planning: From the system point of view, consider the overall architecture design of large-model, multi-model and multi-modal general agent, including hardware architecture, software architecture, data architecture, etc., to meet the requirements of performance, scalability, reliability and security. -Project management and teamwork ability: Have project management and teamwork ability, and be able to lead and coordinate the project team to complete the development and implementation of complex projects. -Problem-solving and communication skills-Problem-solving skills: Excellent independent analysis and problem-solving skills, able to deeply solve various problems existing in the optimization and application of large models. -Communication and presentation skills: able to communicate effectively with team members, cross-departmental colleagues, superiors, etc., including the elaboration of technical solutions, the report of project progress, the feedback of problems, etc. ● Essential knowledge reserve and scientific and technical skills requirements of chief scientist and senior technical expert of industrial intelligence large model multi-model multi-modal general agent: knowledge reserve-deep learning theory: deeply understand neural network architecture, such as Transformer and its variants, master optimization algorithms such as back propagation and gradient descent, and be familiar with concepts such as regularization, over-fitting and under-fitting, and coping methods. -Fundamentals of machine learning: proficient in machine learning paradigms such as supervised learning, unsupervised learning and reinforcement learning, mastering classical algorithms such as clustering, classification and regression, and understanding the methods of model evaluation and selection. -Mathematical foundation: Have solid mathematical knowledge of linear algebra, probability theory, mathematical statistics, calculus, etc., and be able to use mathematical methods to deduce and optimize algorithms. -Knowledge of natural language processing and computer vision: familiar with the techniques of word vector representation, text generation and semantic understanding in natural language processing, as well as the methods of image recognition, object detection and image generation in computer vision. -Computer architecture and parallel computing: Understand the computer

hardware architecture, be familiar with the principle and use of acceleration devices such as GPU and TPU, and master parallel computing and distributed computing technologies, such as multithreading programming and distributed deep learning framework. Science and technology skills-model development and optimization: able to design, develop and optimize large models, including model architecture innovation, parameter adjustment, model compression and quantification, etc., to improve model performance and efficiency. -Multi-modal data processing: master the technology of multi-modal data fusion, representation and processing, and can effectively combine multi-modal data such as text, image and voice for model training and reasoning. -Algorithm innovation and research: innovative, able to carry out cutting-edge algorithm research, propose new model architecture, training methods or optimization strategies, and promote the development of artificial intelligence technology. -Code realization and engineering: proficient in Python, C++ and other programming languages, able to use deep learning frameworks, such as PyTorch, TensorFlow, etc., and have the engineering ability to transform research results into actual products. -Team leadership and collaboration: As the chief scientist, he needs to have the ability to lead and manage the team, guide and train team members, promote team collaboration and promote the smooth progress of the project. ● Necessary knowledge reserves and skill requirements required by chief scientists and senior technical experts in the field of artificial intelligence large model, multi-model and multi-modal general agent, Combing with the technical development trend and industry practice requirements, the system is sorted out:-# # * * 1. Core knowledge reserve **1. ** Basic theory of large model * * * * Deep learning architecture * *: Proficient in model architecture principles such as Transformer, MoE (Mixed Expert) and CoT, and master multimodal alignment (such as QWEN 2.5-). -* * Training and optimization methods * *: Familiar with distributed training, high-efficiency fine-tuning of parameters (PEFT), reinforcement learning (RLHF) and other technologies, and can solve problems such as model illusion and long tail data deviation. -* * Model generalization ability * *: Understand cross-modal reasoning and zero/small sample learning mechanism, such as GPT-4o multimodal dynamic reasoning ability of OpenAI. 2. ** Multi-modal fusion technology * *-* * Cross-modal representation learning * *: Master the unified embedding methods of text, image, voice, video and other modes, such as the Thinker-Talker dual-core architecture of Qwen2.5-Omni to realize real-time synchronous analysis of video and voice. -* * Time-space alignment technology * *: Be familiar with timeline alignment coding (such as TMRoPE) and multimodal data synchronization strategy, and solve scene problems such as audio-video synchronization. -* * Complex document processing * *: Ability to analyze multimodal documents (including tables and charts), which needs to be combined with OCR, layout understanding and semantic association technology. 3. ** Agent system design * *-* * Autonomous decision-making framework * *: Master MetaGPT (Role Cooperative Agent), AutoGen (Dialogue Driven Agent) and other frameworks to realize task planning, tool calling and dynamic environment adaptation. -* * Multi-Agent collaboration * *: Be familiar with the application of federated learning and game theory in multi-agent systems, and optimize the efficiency of

task decomposition and distributed execution. -* * Intelligent integration with body * *: Understand the interaction mechanism between robot action generation (such as the universal 3D modal model of Galaxy) and the physical world, and solve the problems of "limited data" and "action delay". -# # * * II. Key technical skills **1. ** Data engineering ability * *-* * High-quality data construction * *: proficient in multi-modal data cleaning, labeling and enhancement, and need to meet the labeling requirements of millions of kilometers of road test data such as automatic driving. -* * Synthetic data generation * *: A simulation platform (such as Open6DOR) is used to generate large-scale training data with action tags to improve the generalization ability of the robot. -* * Privacy and Compliance Management * *: Be familiar with differential privacy and federated learning technology to ensure data sharing compliance (refer to the National Data Bureau's Data Infrastructure Interconnection Specification). 2. ** System design and optimization * *-* * End-to-end architecture design * *: It can build a "big model+knowledge base+Agent" system (such as Tencent agent development platform), and integrate RAG retrieval, workflow engine and multi-agent collaboration. -* * Optimization of computing power efficiency * *: Grasp the computing power requirements of hybrid cloud deployment, model compression (such as MoE sparse activation), and adaptation of end-side (RTX 3090) and cloud (H100 cluster). -* * Real-time and robustness * *: Optimize the model reasoning delay (such as 300ms speech generation of Qwen2.5-Omni) and design a fault-tolerant mechanism to deal with complex environment fluctuations. 3. ** Exploration of cutting-edge technology * *-* * Breakthrough of embodied intelligence * *: Research on 3D visual point cloud processing and simulation training (such as the hierarchical system of Galaxy universal robot) to improve the success rate of open instruction operation to 95%. -* * Dynamic knowledge management * *: Develop memory-driven RAG (such as Zhiyuan Memo RAG), and realize lifelong learning and personalized service by combining KV cache. -* * Alignment between ethics and safety * *: Design a value alignment mechanism (such as RLHF) to prevent the risk of AI abuse, which meets the requirements of Interim Measures for the Management of Generative AI Services. -# # * * III. Soft ability and industry vision **1. ** Industry-University-Research's integration ability * *-* * Insight into demand scenarios * *: Deeply understand the pain points of vertical industries (such as multi-modal integration for medical diagnosis and dynamic knowledge update for financial risk control), and promote technology adaptation to real business scenarios. -* * Cross-disciplinary cooperation * *: Leading the cooperation between academia and industry, and shortening the technology transformation cycle. 2. ** Strategy and leadership * *-* * Technical route planning * *: Develop a "indomitable spirit" strategy to balance frontier exploration and business landing. -* * Team Management and Incubation * *: Establish multidisciplinary teams (algorithms, hardware, products) and cultivate technical backbones with "full stack capability". -# # * * Fourth, the direction of continuous learning **1. ** Industry standard participation * *: Follow the trends of the Artificial Intelligence Standardization Committee of the Ministry of Industry and Information Technology, and lead the formulation of standards such as multimodal interaction and Agent interface. 2. ** Open-source ecological

construction * *: Contribute and maintain open-source projects (such as HuggingGPT and mixed-element multimodal model) and promote the popularization of technology. 3. ** Vision of international competition * *: Pay attention to the technical differences between China, the United States and Europe (such as the cooperation between OpenAI and Figure AI), and lay out patents and core technical barriers. -# # * * Typical competency benchmarking * *-* Tencent mixed-element team * *: It needs to have the ability of large-scale model research and development (TurboS pedestal optimization), Agent platform construction (zero-code multi-agent collaboration) and knowledge base integration. -* * Galaxy Universal Robot * *: Full-link technology control that requires three-dimensional visual modeling, simulation data generation and large-scale model scheduling. -* * Academic leader * *: It is necessary to be able to connect theory (thinking chain reasoning) with application (industrial brain) and promote the paradigm innovation of "AI+ industry". -* * Summary * *: Chief scientists and senior technical experts need to build a trinity capability system of "technical depth+industry breadth+strategic height", which not only overcomes technical difficulties such as multi-modal alignment and Agent independent planning, but also promotes the integration and standardization of Industry-University-Research. Finally,